

Virtual Energy Storage from TCLs using QoS persevering local randomized control

Austin R. Coffman*
University of Florida
Gainesville, FL
bubbaroney@ufl.edu

Ana Bušić
INRIA
Paris, France
ana.busic@inria.fr

Prabir Barooah
University of Florida
Gainesville, FL
pbarooah@ufl.edu

ABSTRACT

We propose a control architecture for distributed coordination of a collection of on/off TCLs (thermostatically controlled loads), such as residential air conditioners, to provide the same service to the power grid as a large battery. This involves a collection of loads to coordinate their on/off decisions so that the aggregate power consumption profile tracks a grid-supplied reference. A key constraint is to maintain each consumer's quality of service (QoS). Recent works have proposed randomization at the loads. Thermostats at the loads are replaced by a randomized controller, and the grid broadcasts a scalar to all loads, which tunes the probability of turning on or off at each load depending on its state. In this paper we propose a modification of a previous design by Meyn and Bušić. The previous design by Meyn and Bušić ensures that the indoor temperature remains within a pre-specified bound, but other QoS metrics, especially the frequency of turning on and off was not limited. The controller we propose can be tuned to reduce the cycling rate of a TCL to any desired degree. The proposed design is compared against the design by Meyn and Bušić and another well cited design in the literature on control of TCL populations, by Mathieu *et al.* We show through simulations that the proposed controller is able to reduce the cycling of individual ACs compared to the previous designs with little loss in tracking of the grid-supplied reference signal.

CCS CONCEPTS

• **Hardware** → **Smart grid; Power estimation and optimization;**

KEYWORDS

Virtual Energy Storage, Distributed Control, Demand Response, Randomized Control

ACM Reference Format:

Austin R. Coffman, Ana Bušić, and Prabir Barooah. 2018. Virtual Energy Storage from TCLs using QoS persevering local randomized control. In *The 5th ACM International Conference on Systems for Built Environments*

*This is the corresponding author

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

BuildSys '18, November 7–8, 2018, Shenzhen, China

© 2018 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-1-4503-5951-1/18/11...\$15.00
<https://doi.org/10.1145/3276774.3276777>

(BuildSys '18), November 7–8, 2018, Shenzhen, China. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3276774.3276777>

1 INTRODUCTION

Reliable operation of the power grid requires balancing demand and supply of power at all time scales. The time variation and unpredictability of renewable energy sources such as solar and wind make it challenging. Apart from expensive batteries, a complementary and inexpensive possibility is to harness the inherent flexibility in demand of many types of loads. Loads have been used for Demand Response (DR) for a long time, which is traditionally meant as a temporary reduction in demand to help the grid. It is recognized now that loads can supply a range of services to the grid beyond DR [1, 14]. With appropriate control, loads can vary their demand up and down around a baseline so that the deviation from the baseline appears like the charging and discharging of a battery to the grid. We call this Virtual Energy Storage (VES) from smart loads. A key constraint in using loads to provide any kind of grid-support service is that consumer's quality of service (QoS) must not be compromised.

The topic of this paper is design of the control architecture for a collection of thermostatically controlled loads (TCLs) to provide VES while maintaining strict bounds on the consumers' QoS. TCLs include residential air conditioners (ACs), water heaters, refrigerators etc. We focus on residential ACs that are on/off type; meaning their power consumption can only take two values: zero and a positive constant. For ACs, there are at least two primary measures of consumer's QoS: indoor temperature and cycling frequency. Cycling frequency refers to the number of times a load turns on and off in a given period. Short-cycling, which refers to frequent turning on and off, is to be avoided since that can damage equipment.

The problem is challenging on many counts. First is computational complexity. A thousand on/off loads means at any instant there are 2^{10^3} possible decisions. Due to QoS constraints that make current decisions dependent on the past, the decision space is even bigger. Second, since loads are distributed geographically over a large area and QoS constraints are local, the control architecture must be non-centralized. However, loads' actions must be coordinated so that together they deliver the service that the balancing authority (BA) asks for.

There is a large and growing literature on coordination of collections of TCLs. Several distinct heuristics have been developed to address computational complexity. One control approach is for the BA to broadcast a common thermostat set point change [2, 5, 10, 18]. The control signal is decided based on a model with set point change as input and aggregate power consumption as output. A limitation of this approach is that the control relies on extremely small

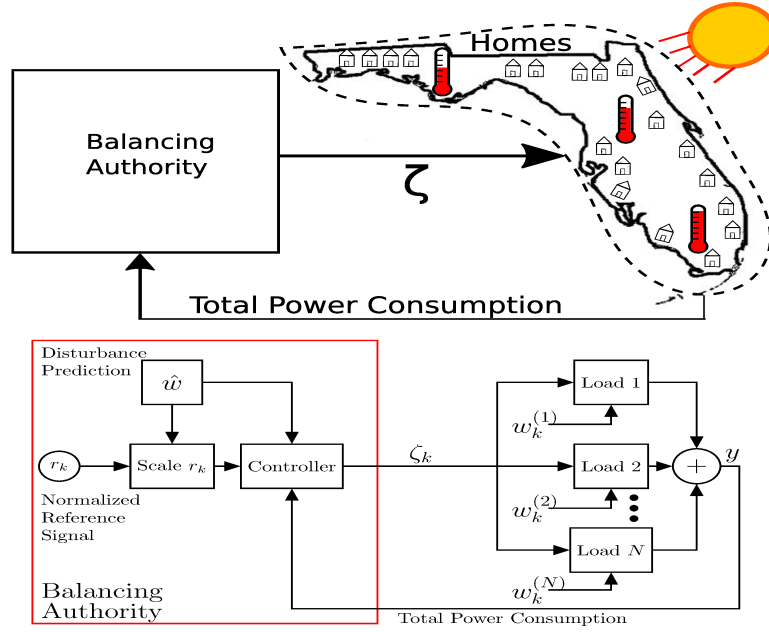


Figure 1: (Top): High level architectural diagram of information requirements: ζ is a scalar broadcast to all loads. (Bottom): Control architecture.

set point change command, as little as 0.0025 °C. Most residential thermostats have a setpoint resolution of 0.5 °C or more, and the deadband of a thermostat is also much wider than 0.0025 °C. Large transient oscillations due to synchronization of loads is another limitation [2]. The work [18] offered a partial amelioration of the resolution issue by dividing the loads into clusters and sending distinct, but coarse set point changes, to every cluster.

Another widely used approach is based on first modeling the distribution of temperatures of the population as a finite-state Markov chain model by dividing temperature into bins; see [13, 16, 20] and references therein. The population density model is developed assuming that the local control algorithm that operates the load remains in place. We refer to this local controller as the “thermostat controller”. The thermostat controller ensures that the indoor temperature remains within a deadband. Control by the BA therefore involves computing quantities such as the fraction of loads that should turn on or off. Converting that global command to individual on/off decisions to ensure that the collection of TCLs track a desired total power consumption reference, while at the same time satisfying each loads’ local QoS constraints, is a challenging problem. Although various heuristics are developed to this end (see e.g., [16, 20]), these challenges are not addressed in a principled manner.

The control architecture we adopt in this paper is taken from a recent series of papers [3, 4, 7, 8, 17]. This architecture (Figure 1) involves intelligence both at the BA (global) and the load (local). The local intelligence at the load is a randomized controller that is meant to replace the thermostat, but is designed so that the resulting behavior - from the point of view of the consumer - is indistinguishable from that of a thermostat. The controller at the

load has a tunable parameter ζ that determines the degree of randomness in this decision making. Higher ζ (positive) increases the probability of turning on while lower ζ (negative) increases the probability of turning off. For $\zeta = 0$ operation is similar to thermostat control. The controller at the BA computes ζ and broadcasts to every load. Randomization at the loads helps avoid synchronization, and simplifies the control computation at the BA. The latter is crucial: randomization enables approximating the complex high-dimensional non-linear dynamics of the collection by a low-order input-output system with ζ being the input signal and total power consumption of the collection as the output. The BA can therefore use a simple dynamic compensator to compute the control command ζ based on measurement of total power consumption of the TCLs.

The contribution of this paper is a novel design of the local randomized controller. The local randomized control at the loads used here is inspired from [4, 7]. The key innovation over [4, 7] is that while the controllers in those references only ensured one QoS metric, namely, indoor temperature of an AC, the controller proposed here ensures that in addition, the cycling QoS metric is also maintained. In particular, the local randomized controller proposed here ensures that the number of cycles (on/off transitions) that an AC undergoes in a given time interval does not vary much from its baseline behavior (no interference from the BA). The proposed local controller is called CARC (Cycling Aware Randomized Control). The control command at the BA, ζ , is computed by using a PI (proportional-integral) control scheme.

We provide simulation evaluations of the proposed controller with that in [4, 7] and that in [16]. Simulations show that the proposed controller is able to improve the cycling QoS over both of

these competing designs, with little or no loss in performance in tracking the BA supplied reference.

2 PROBLEM FORMULATION: NEEDS OF THE GRID

The BA needs resources to provide the *net load*, \tilde{d}_k , (where k denotes the discrete time index) which is the difference between the baseline demand and renewable generation. The baseline demand is the demand when the loads are not providing any grid-support service, i.e., there is no interference from the grid on power demand of the loads.

The slowly-varying component of the net load is ideally provided by conventional generators that have limited ramping ability. This component, denoted by \tilde{d}_k^{LP} can be obtained by low-pass filtering the net-load. The remainder is denoted by $\tilde{d}_k^{HP} = \tilde{d}_k - \tilde{d}_k^{LP}$. This "high-pass" component is zero-mean, and can be provided by VES and actual energy storage. A variant of this methodology is used by many BAs today; the \tilde{d}_k^{HP} is provided by fast moving generators, flywheels, and batteries.

In BPA (Bonneville Power Administration), a BA in the pacific northwest with a high share of wind penetration, the *balancing reserve* is close to the signal \tilde{d}_k^{HP} . An example of BPA's balancing reserve signal is shown in Figure 2. The data is obtained from <http://tinyurl.com/ybqlq8w>.

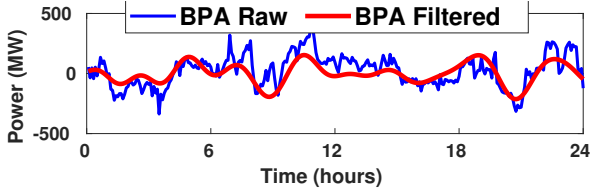


Figure 2: Illustration of storage requirement \tilde{d}_k^{HP} (BPA's balancing reserve). The reference r_k is obtained by filtering the balancing reserve signal.

Loads providing VES can be characterized by the frequency of demand variation they can tolerate without violating QoS constraints [9, 11, 12, 17]. Therefore, an appropriate reference signal for the collection of ACs to track, denoted by $\tilde{P}^{ref}(z)$, can be obtained by bandpass filtering the signal $\tilde{P}^{ref}(z) = F_{BP}(z)\tilde{d}_k^{HP}$, where $F_{BP}(z)$ is a bandpass filter, $\tilde{P}^{ref}(z)$ is the reference signal measured in Watt (or kW, MW, GW etc.), and z is the z -transform variable. The passband of $F_{BP}(z)$ must be chosen according to the ability of the aggregate. To determine the appropriate passband, the BA needs knowledge of the frequency response of the collection of loads. We describe in Section 4.2.1 how frequency response of the collection of ACs, can be identified, with ζ as input and total power consumption deviation (from the baseline) as output. The passband is then chosen to be range of frequencies in which the aggregate model has high gain and small phase lag. An example of a filtered signal, $\tilde{P}^{ref}(z)$, with passband appropriate for a collection of AC's is shown in Figure 2.

Let the power consumption of the collection of loads be denoted by P_k , and P_k^* be the value of P_k in the baseline scenario (no interference from the BA). The deviation of the total power consumption of all the loads is $\tilde{P}_k := P_k - P_k^*$. The control problem for the BA is a tracking problem: \tilde{P}_k should track the reference \tilde{P}_k^{ref} . Furthermore, the BA should only provide a reference that the loads *can* track. The frequency content of the reference signal \tilde{P}_k^{ref} is only a necessary condition for successful operation. The collection of loads additionally have power and energy constraints, like a battery, that need to be respected. Reference signals that exceed the value of power consumption when all loads are on (power constraint), or require to fast a consumption rate (energy constraint) will result in poor tracking or violation of individual ACs QoS.

3 LOCAL CONTROL AT THE AC'S, AND AGGREGATE BEHAVIOR

3.1 Dynamic model of indoor temperature

Consider an AC with p^{rated} being the rated electrical power consumption. A simple model of the indoor temperature θ is, $\dot{\theta}(t) = -\frac{1}{RC}\theta + \frac{q_0}{C}q_{ac}(t) + w(t)$, where $w(t) := \frac{1}{RC}\theta_a(t) + \frac{1}{C}q_{int}(t)$ is a time varying disturbance with θ_a the ambient temperature and q_{int} the internal cooling load, R is the resistance to heat flow offered by the building structure and C is the thermal capacitance of the building. The term q_{int} captures both occupant induced heat gain and solar heat gain. The $q_{ac}(t)$ is the heat injected into the building by the AC. The control signal $u(t)$ is binary: it can be either 1 (on) or 0 (off). Denoting by $q_0 := COP^{rated}$ the rated thermal power consumption of the AC, COP being its coefficient of performance, $q_{ac}(t) = -q_0$ if $u(t) = 1$ and $q_{ac}(t) = 0$ if $u(t) = 0$. We now have the binary discrete control signal u_k : 1 when on and 0 when off. The power consumption of the AC (in kW) is $p^{rated}u_k$.

3.2 Deterministic control (Thermostat)

The control logic in a thermostat that operates a residential AC is usually based on a deadband (δ) around a user-specified set point (θ_{set}). The AC is turned on if the measured indoor temperature exceeds the upper limit $\theta_{max} = \theta_{set} + \delta$ and turned off if the temperature drops below the lower limit $\theta_{min} = \theta_{set} - \delta$. When $\theta(t)$ is between θ_{min} and θ_{max} , the previous decision (on or off) is maintained. This control strategy is *deterministic*: the on/off status of the AC at $k+1$ is a function of the temperature and on/off status at k : $u_{k+1} = u_{k+1}(\theta_k, u_k)$.

3.3 Randomized control basics

In randomized control [4, 7], the mapping $u(\theta_k, u_k) \mapsto u_{k+1}$ is no longer deterministic; it takes the two possible values (0 and 1) with certain probabilities. First define a state space for the AC. Every element x of the state space X is a pair, $x = (u, \theta)$, where $u \in \{0, 1\}$ is the "mode" (either on or off) and $\theta \in \mathbb{R}$ is the temperature of the house. For a state $x \in X$, we denote by x^u the mode and x^θ the temperature θ . In the sequel, we use \oplus and \ominus interchangeably with 1 and 0, respectively, as it is more intuitive for on and off.

A randomized controller is a rule to determine the probabilities of x_{k+1}^u being 0 or 1 given the knowledge up to the current time,

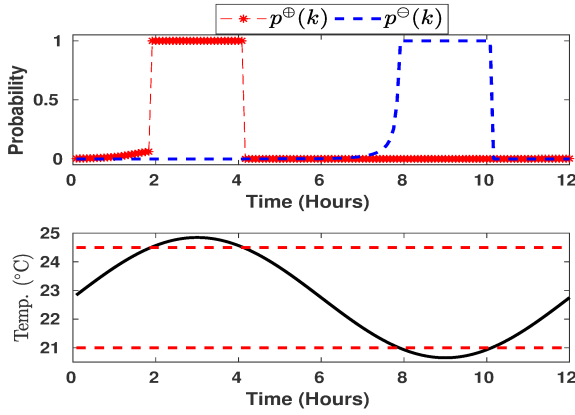


Figure 3: Switching probabilities as a function of an arbitrarily chosen indoor temperature trajectory (bottom), for the baseline case $\zeta = 0$.

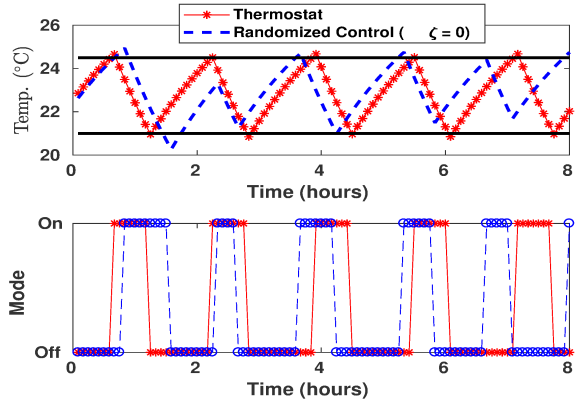


Figure 4: Comparison of thermostat control and baseline randomized control with $\zeta = 0$. The indoor temperature occasionally exceeds the deadband, which occurs due to thermal inertia of the air and the delay caused by discrete time control.

k . The controller is therefore completely specified by a Markov operator R_ζ . The quantity $R_\zeta(x, y^u)$ is the probability of the mode being y^u at the next instant $k + 1$ given the state is x at the current time k . The Markov operator is parameterized by a real number ζ . Designing a controller is equivalent to designing the operator $R_\zeta(x, y)$ for all $x, y \in X$ and all $\zeta \in \mathbb{R}$.

The evolution of the temperature is governed by the dynamic system described in section 3.1. We can represent it as another Markov operator $Q_w(x, y^\theta)$: the probability of the temperature being y^θ at the next instant $k + 1$ given the state is x at the current time k . Since temperature evolution is a deterministic dynamic system, the transition probabilities $Q_w(x, y^\theta)$ will be Dirac-delta functions. The next state will depend on the disturbance w , hence the operator $Q_w(x, y^\theta)$ is parameterized by the disturbance w .

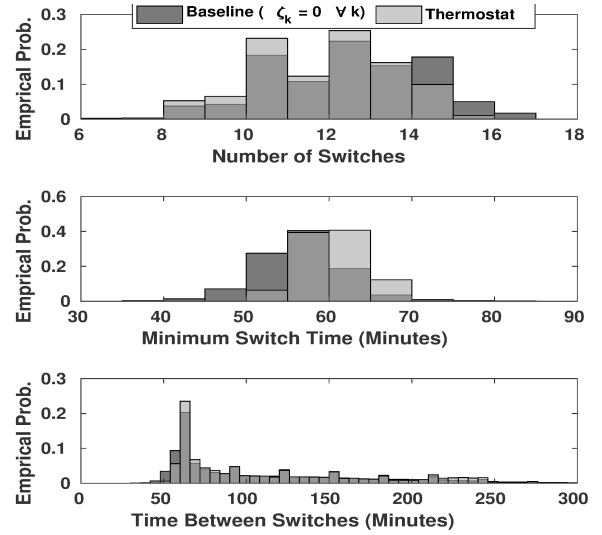


Figure 5: Empirical probabilities comparing our proposed nominal design against the deterministic thermostat controller.

Combining the effect of the control and the disturbance, we get the total probability of transitioning from state x to y over one time period, which is denoted by $P_{\zeta, w}(x, y)$:

$$P_{\zeta, w}(x, y) = R_\zeta(x, y^u)Q_w(x, y^\theta), \quad x, y, \in X \quad (1)$$

The model (1) is an extension of the finite-state Markov model of a TCL in [4]; here the state space X is infinite [3].

The controller $R_\zeta(x, y^u)$ has to be designed in such a way that the following properties are satisfied. One, the operator $R_0(x, y^u)$ in the *baseline* case $\zeta = 0$ (corresponding to no interference from the grid) mimics the behavior of the deterministic control. Meaning, it has to turn the AC on if indoor temperature exceeds θ_{\max} and off if the indoor temperature falls below θ_{\min} , and avoid short cycling. Two, if ζ is positive, the probability of turning on should increase and if ζ is negative, probability of turning off should increase. In effect, the variable ζ can be used by the BA as the control signal to manipulate the power consumption of an AC.

Simulation results showing both of these requirements are presented in Figure 3-4 for a single AC and in Figure 5 for a collection of ACs. Parameters of the detailed simulation are described later. The baseline scenario is seemingly identical to the thermostat controller, except that it causes a AC to occasionally switch on or off more than the thermostat (Figure 5 Top).

3.4 Randomized control of MB-RC

To construct the baseline control law, $R_0(x, y^u)$, two switching probability functions need to be defined. To construct these functions, first consider the two random variables:

θ^+ : the temperature above which the A/C turns on (after having started in the off state, or since the last off state)

θ^- : the temperature below which the A/C turns off (after having started in the on state, or since the last on state).

Consider the usual heating scenario, when the AC has been off and temperature is increasing. Then, at time k , the probability of switching on is

$$\begin{aligned} p^{\oplus}(k) &:= P(X_k^u = 1 | X_{k-1}^u = 0, X_{k-2}^u = 0, \dots) \\ &= P(\theta^+ \leq \theta_k | \theta^+ > \theta_{k-1}) \\ &= \frac{P(\theta_{k-1} < \theta^+ \leq \theta_k)}{P(\theta^+ > \theta_{k-1})} \\ &= \frac{F^{\oplus}(\theta_k) - F^{\oplus}(\theta_{k-1})}{1 - F^{\oplus}(\theta_{k-1})} \end{aligned}$$

Additionally for a cooling scenario, when the AC has been on and the temperature is decreasing. Then, at time k , the probability of switching off is

$$\begin{aligned} p^{\ominus}(k) &:= P(X_k^u = 0 | X_{k-1}^u = 1, X_{k-2}^u = 1, \dots) \\ &= P(\theta^- \geq \theta_k | \theta^- \leq \theta_{k-1}) \\ &= \frac{P(\theta_k < \theta^- \leq \theta_{k-1})}{P(\theta^- \leq \theta_{k-1})} \\ &= \frac{F^{\ominus}(\theta_{k-1}) - F^{\ominus}(\theta_k)}{F^{\ominus}(\theta_{k-1})} \end{aligned}$$

where F^{\oplus}/F^{\ominus} is the CDF of the r.v. θ^+/θ^- . Notice that in both scenarios we have assumed (third line) that temperature has been increasing while off (decreasing while on), so that $\theta_{k-1} < \theta_k$ ($\theta_k < \theta_{k-1}$). To ensure positivity of probabilities under all scenarios, e.g., if temperature decreases (increases) when AC is off (on), we use

$$\begin{aligned} p^{\oplus}(k) &= \frac{[F^{\oplus}(\theta_k) - F^{\oplus}(\theta_{k-1})]_+}{1 - F^{\oplus}(\theta_{k-1})} \\ p^{\ominus}(k) &= \frac{[F^{\ominus}(\theta_{k-1}) - F^{\ominus}(\theta_k)]_+}{F^{\ominus}(\theta_{k-1})}. \end{aligned} \quad (2)$$

where $[x]_+ = \max(x, 0)$. The two CDF's appearing in (2) are additional design choices and are taken from [4] as,

$$\begin{aligned} F^{\oplus}(z) &= \exp(-(\theta_{\max} - z)^{\rho} / (2\sigma^{\rho})) \\ F^{\ominus}(z) &= 1 - F^{\oplus}(\theta_{\max} + \theta_{\min} - z) \end{aligned} \quad (3)$$

where ρ and σ are design parameters. In this paper we use $\rho = 0.75$ and $\sigma = 0.02$.

With the previous developments the baseline (nominal) control law, $R_0(x, y^u)$, can be stated explicitly as:

$$\begin{aligned} R_0(x, 1) &= P(X_k^u = 1 | X_{k-1} = x) \\ &= \begin{cases} p^{\oplus}(x^{\theta}), & x^u = \ominus \\ 1 - p^{\oplus}(x^{\theta}), & x^u = \oplus \end{cases} \\ R_0(x, 0) &= P(X_k^u = 0 | X_{k-1} = x) \\ &= \begin{cases} p^{\ominus}(x^{\theta}), & x^u = \oplus \\ 1 - p^{\ominus}(x^{\theta}), & x^u = \ominus \end{cases} \end{aligned} \quad (4)$$

With abuse of notation (state augmentation is required to illustrate the 1st order Markov property) the control law (when $\zeta \neq 0$), proposed in [4, 7], is the *myopic policy*:

$$R_{\zeta}(x, y^u) = R_0(x, y^u) \exp(\zeta U(y^u) - \Lambda_{\zeta}(x)) \quad (6)$$

where $U(\cdot)$ is the *utility function*:

$$U(x) = U(x^u) = \begin{cases} 1 & x^u = \oplus \\ 0 & x^u = \ominus \end{cases} \quad (7)$$

and $\Lambda_{\zeta}(x)$ is the normalization constant to make the probabilities sum to 1.

This translates to switching probability functions for $\zeta \neq 0$:

$$\begin{aligned} p_{\zeta}^{\oplus}(k) &= \frac{p^{\oplus}(k) \exp(\zeta)}{\exp(\Lambda_{\zeta}(x))} \\ p_{\zeta}^{\ominus}(k) &= \frac{p^{\ominus}(k)}{\exp(\Lambda_{\zeta}(x))} \end{aligned} \quad (8)$$

The developments of this section are referred to as the Meyn & Bušić - Randomized Control (MB-RC) [4, 7]. Notice that the controller uses info from the past two time instances, so to represent it as a Markov operator we need to augment the state. We avoid doing so here to prevent the notation from becoming too cumbersome.

3.5 Proposed local controller: CARC

In this paper we propose a change to the local randomized controller (6) to ensure that the AC unit does not cycle excessively. We coin this new design as the Cycling Aware Randomized Control (CARC) Algorithm. To describe the CARC algorithm, we first introduce a function $f(\cdot) : \{x^u\}_{k-\tau-1}^k \mapsto \mathbb{R}^+$, which is defined as:

$$f(\{x^u\}_{k-\tau-1}^k) = \exp\left(-\alpha \sum_{t=0}^{\tau} |x_{k-t}^{(i),u} - x_{k-t-1}^{(i),u}|\right) \quad (9)$$

for $k > \tau$, where τ is a horizon window and α is a decay rate. Both τ and α are design variables. This function, $f(\cdot)$, multiplicatively effects ζ leading to the CARC randomized local control law:

$$R_{\zeta}(x, y^u) = R_0(x, y^u) \exp(\zeta f(\{x^u\}_{k-\tau-1}^k) U(y^u) - \bar{\Lambda}_{\zeta}(x)), \quad (10)$$

where $R_0(x, y^u)$ and $U(y^u)$ are the same as defined previously for the MB-RC algorithm, and $\bar{\Lambda}_{\zeta}(x)$ is a normalizing constant.

Both CARC and MB-RC share the same nominal design (when $\zeta = 0$). The CARC algorithm effectively reduces the magnitude of ζ locally if it detects that the AC has already cycled quite a few times, and thus avoids excessive cycling. This is achieved through the design of the decaying exponential in (9). If an AC unit begins to cycle excessively, the summand in (9) will be large and in turn $f(\cdot)$ will be small, and so the product $f(\cdot)\zeta$ will be small as well. This causes the unit to behave close to nominal operation; as if the ζ broadcast by the BA were 0. If the AC has not cycled many times within τ then the summand in (9) will be zero, $f(\cdot)$ will be one, and the scaling will have no effect.

The success of the CARC algorithm in ensuring that cycling frequency does not increase much more than what the thermostat controller would do without any interference from the grid is an appropriate nominal (baseline) design. The baseline randomized controller (when $\zeta = 0$) in our design behaves similarly to a thermostat (see Figures 3- 5) in Section 3.3).

To execute the proposed control algorithm at an AC, it needs to be able to receive the broadcast ζ_k from the BA. It also needs its on/off state and the indoor temperature, which are local measurements. No other information is needed. The computations involved are quite

simple and can be performed with the help of a random number generator and a look up table of the switching probabilities.

3.6 Aggregate behavior of N ACs with local randomized control

Let μ_k be the pdf (defined over the state space X) of the state of an AC when it operates according to the randomized control law described in Section 3.4 or Section 3.5. The pdf (left) and the relevant output, the probability, at time k , of the state being in the “on” mode, (right) evolve according to

$$\mu_{k+1} = \mu_k P_{\zeta_k, w_k} \quad \gamma_k := P(x_k^u = \oplus) = \int_{R^\oplus} \mu_k dx, \quad (11)$$

As a system with inputs ζ, w and output γ , (11) is an infinite dimensional dynamic system that is linear in the state but nonlinear in the inputs. To make the connection between the individual and the aggregate, consider N homogeneous loads operating under the same randomized policy and subject to the same inputs (ζ_k, w_k) . Define the fraction of loads that are on at k :

$$y_k^{(N)} := \frac{1}{N} \sum_{i=1}^N U(x_k^{(i)}) \quad (12)$$

We assume a mean-field limit holds (proven for a similar setting in [17]). That is, if N homogeneous loads are subjected to the same inputs ζ and w , as $N \rightarrow \infty$ the fraction of loads that are on at time k approaches the probability of a single load being on at k :

$$\lim_{N \rightarrow \infty} y_k^{(N)} = \gamma_k \quad (13)$$

In the sequel, for simplicity we drop the superscript N from all quantities that contain an average over N .

4 DECISION MAKING AT THE BA

The control problem faced by the BA was discussed in Section 2. We now pose the problem more precisely and connect it to the aggregate behavior described above.

4.1 The tracking problem

When $\zeta_k \equiv 0$ and $w_k \equiv w$, then γ_k reaches a constant steady state value as $k \rightarrow \infty$, which we denote by $\gamma^*(w)$. In other words, $\gamma^*(w)$ is the fraction of the loads that are on at steady state for a constant disturbance w under baseline randomized control. We now define the deviation signal $\tilde{\gamma}_k := \gamma_k - \gamma^*(w_k)$.

When the system is operating in baseline conditions, i.e., $\zeta_k \equiv 0$, but with a time-varying disturbance, then $\gamma_k = \gamma^*(w_k)$ at every k if the response of the aggregate to the disturbance is instantaneous. Assuming such a speedy response, we have $\tilde{\gamma}_k \equiv 0$ for baseline operation. When $\zeta_k \equiv \zeta > 0$ then $\tilde{\gamma}_k$ converges to a positive value that is larger for larger ζ . This follows from the design of the controlled transition probability operator R_ζ that was described in Section 3.3. Similarly, if $\zeta_k \equiv \zeta < 0$ then $\tilde{\gamma}_k$ converges to a negative value that is smaller for smaller ζ .

Recall y_k defined in (12); the empirical counterpart to γ_k . We similarly define the empirical counterparts to $\gamma^*(w)$ and $\tilde{\gamma}_k$, and call them $y^*(w)$ and \tilde{y}_k . Therefore,

$$\tilde{y} = y_k - y_k^*. \quad (14)$$

Assuming the mean field limit holds, $y^*(w) \rightarrow \gamma^*(w)$ and $\tilde{y}_k \rightarrow \tilde{\gamma}_k$ as $N \rightarrow \infty$.

Recall the tracking objective from Section 2: the power consumption deviation of the collection \tilde{P}_k should track an exogenous reference signal \tilde{P}_k^{ref} . When the loads are homogeneous, $\tilde{P}_k = P_k - P_k^* = N p^{\text{rated}}(y_k - y_k^*) = N p^{\text{rated}} \tilde{y}_k$. The problem is then equivalent to \tilde{y}_k tracking the *normalized reference* $r_k := \frac{1}{N p^{\text{rated}}} \tilde{P}_k^{ref}$. Because of the mean field limit, \tilde{y}_k tracking r_k is (approximately) equivalent to $\tilde{\gamma}_k$ tracking r_k . For the BA, the control design problem can be posed in terms of a single load: choose the control command ζ_k so that the deviation probability $\tilde{\gamma}_k$ of a load tracks the normalized reference r_k .

4.2 Computing r_k , the reference for the collection

4.2.1 Frequency response of a collection of ACs. In order to obtain the passband for the bandpass filter described in Section 2, the frequency response of the collection of AC's with ζ_k as the input and \tilde{y}_k as the output, $G_{\zeta \tilde{y}}(e^{j\Omega})$, needs to be computed. This frequency response represents the linear approximation to the non-linear mean field model around an equilibrium point, and because of the mean field limit the approximation will predict the change in the fraction of ACs on, \tilde{y}_k , over baseline, y_k^* , due to ζ_k .

In order to estimate the frequency response, the collection of AC's is simulated with a constant disturbance, w^* , and a zero mean ζ . This corresponds to estimating a model about the equilibrium point ($\zeta = 0, w = w^*$). The corresponding normalized power deviation, $\tilde{\gamma}_k$, is recorded and used along with the elected inputs to compute the Empirical transfer function estimate (ETFE). The ETFE can be used to fit a discrete time transfer function. The results of this estimation procedure, for loads with the CARC algorithm and with MB-RC, are shown in Figure 8. It can be observed that the frequency response of the aggregate is not effected by the introduction of the local QoS constraint.

4.2.2 Computing the Baseline power P_k^* . Recall, the tracking objective for VES is for a collection of AC's to track an exogenous reference signal (\tilde{P}_k^{ref}) that represents power deviation from baseline consumption (P_k^*). Baseline consumption represents the power consumption that would have occurred under no VES. In order to compute the baseline power consumption, the collection of AC's are simulated, each equipped with the local randomized controller, with $\zeta_k = 0 \forall k$, and the exogenous disturbance trajectory, which can be obtained from weather forecasts. This requires a prediction of the disturbance for the desired time horizon. We assume the ability to predict the disturbance accurately 1 day ahead, which is reasonable as some studies have done predictions months in advanced [19], and BA's frequently use such demand forecasts. The baseline consumption for the control experiments is shown in Figure 6.

4.2.3 Computing the Capacity c_k . Under time varying weather conditions, the magnitude of trackable references is also time varying. This is illustrated in Figure 6, observe at 6 AM only 20% of the AC units are on. Clearly it is impossible to turn off more than 20% additional units, and any reference that has a value in excess of

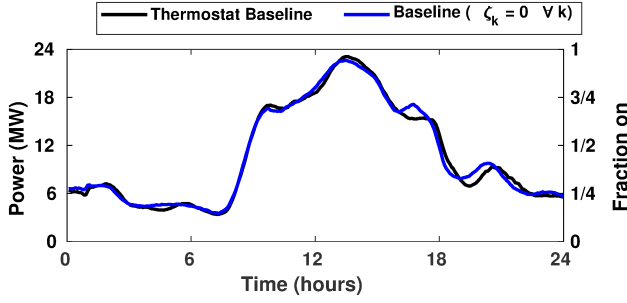


Figure 6: Baseline power consumption of the collection of 10000 AC's for the simulation scenarios.

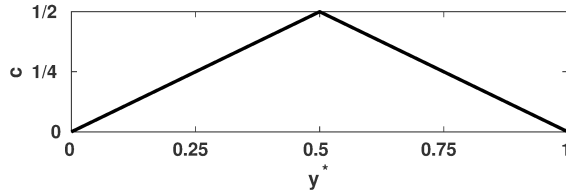


Figure 7: Capacity c_k as a function of normalized baseline power consumption, y^* .

± 0.20 at 6 AM cannot be tracked no matter the control complexity. This is largely dependent on the disturbance $w(t)$, which is dependent on the time varying ambient temperature $\theta_a(t)$. In order to ensure that percentage quantities greater in magnitude than baseline consumption are not requested, a capacity $c_k(y^*)$, as a function of the baseline, is computed and used to scale the reference r_k [8]. The capacity function we use is illustrated in Figure 7.

4.3 Computing the control signal ζ

We propose to use a PI controller at the BA level to compute the control signal ζ_k . Model Predictive Control (MPC) has been used to compute the control signal ζ_k in [8]. However, it was found that this level of sophistication may not be necessary as the aggregate power behaves linearly (Figure 8). In this work, at time k the BA computes the control signal ζ_k using:

$$\zeta_k = K_p(c_k r_k - \tilde{y}_k) + K_I \sum_{i=0}^k c_i r_i - \tilde{y}_i \quad (15)$$

Where K_p and K_I are the proportional and integral gains for the PI controller, respectively.

The normalized baseline power consumption, y_k^* , reference signal, r_k , and capacity, c_k , are computed with the methods described in Sections 2, 4.2.2, and 4.2.3, respectively. We assume that the BA has knowledge of the approximate distribution of parameters of the ACs, and a forecast of the weather for the next 24 hours, so all these signals are computable at the BA by using the methods described previously. To compute the control command ζ_k by using (15), the BA also needs the feedback signal \tilde{y} , which can be obtained if it knows the total power consumption of the aggregate and the

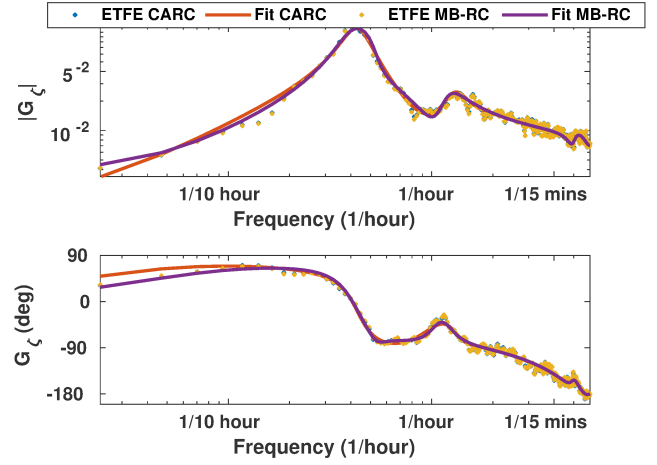


Figure 8: Bode plot of the finite dimensional LTI system that approximates the aggregate model. ETFC CARC/Fit CARC: ETFC and TF Fit of the simulation with CARC algorithm. ETFC MB-RC/Fit MB-RC: ETFC and TF Fit of the simulation with MB-RC.

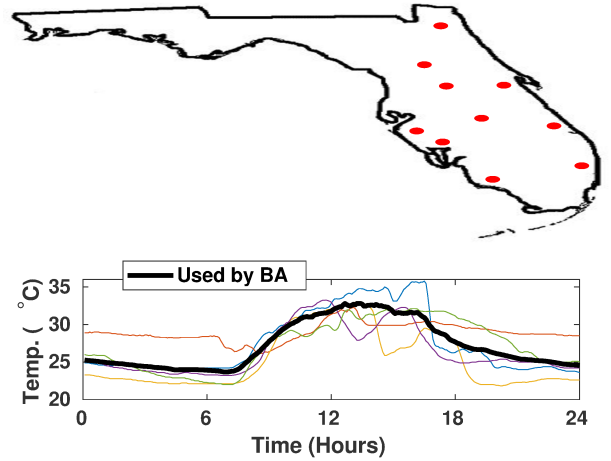


Figure 9: Top: Locations of weather stations (red dots), Bottom: Temperatures seen by five weather stations and temperature prediction used by the BA.

baseline power. In this paper we assume that each AC sends back its on/off state at every sampling period (5 minutes) to the BA. The measurement of the total power consumption is thus known to the BA.

5 PERFORMANCE EVALUATION

5.1 Simulation set up

To evaluate the performance of the proposed method, we compare the QoS results of the new proposed local controller (CARC) with those of [16] and Meyn & Bušić - Randomized Control (MB-RC). We denote the control architecture proposed in [16] as Mathieu Koch

Calloway - Probabilistic Switching (MKC-PS), where in Section 5.2 further details on this method will be provided. The QoS metrics we are interested in are of two types: (i) Local QoS constraints of the individual AC and (ii) Global QoS of the Balancing Authority. The two local QoS measures of interest are indoor temperature and AC cycling frequency. At the BA-level we are interested in tracking of the total power deviation to that of an exogenous reference signal, so the QoS metric is tracking error.

The BA performs all the necessary calculations, including the off-line system identification, with heterogeneous AC parameters (Table 1), but assuming that each AC is subjected to the same weather. This is done to test robustness to uncertainties, since during closed-loop operation each AC is subjected to a different disturbance value. To illustrate the range of the individual disturbance values, 5 sample paths for the ambient temperature are provided in Figure 9. The parameters for the heterogeneous loads are described in Table 1 and are taken within the bounds provided from Table 1 in [15]. These parameters are used in all the closed-loop simulations. Additionally, the parameters specific to the CARC algorithm are also given in Table 1.

Table 1: Parameters

Classification	Symbol	Value	Units
AC parameters			
Resistance for load model	R	$U[2, 2.25]^*$	K/kW
Capacitance for load model	C	$U[2, 2.25]$	kWh/K
Rated elec. power	p^{rated}	$U[4.4, 6.5]$	kW
Coeff. of Perf.	COP	2.5	NA
Temperature set point	θ_{set}	$U[21, 23]$	°C
Temperature deadband	δ	$U[0.75, 1]$	°C
CARC Parameters			
QoS Horizon	τ	9	samples
QoS Decay Rate	α	2	N/A
Parameters for the BA			
Proportional Gain	K_p	27	N/A
Integral Gain	K_I	35	N/A

* $U[a, b]$ is uniformly distributed between a and b.

Geographically, we imagine the BA is located in the state of Florida in the southeastern United States, and its territory covers almost the entire state; see Figure 9. The ambient temperature prediction it uses in all of its calculations is shown as the thick black line in Figure 9. This signal is obtained by averaging the temperature measurements collected from ten weather stations in Florida during two days in June 2017. The stations are shown in Figure 9. The BA predicts the disturbance signal from weather prediction by using the R, C values assumed and assuming that the occupant-induced heat-gain q_{int} is 0 at every load and every instant. We also envision a future for Florida, which is called the Sunshine State, in which it has a high penetration of solar. To mimic such a scenario, the reference signal r_k is chosen by filtering BPA's balancing reserves data as described in Section 2. A Butterworth filter with passband $[\frac{1}{3}, \frac{1}{1.5}]$ (1/hour) was used. The choice of the passband is informed by the identified model, since the LTI model observed high gain in this region (see Figure 8).

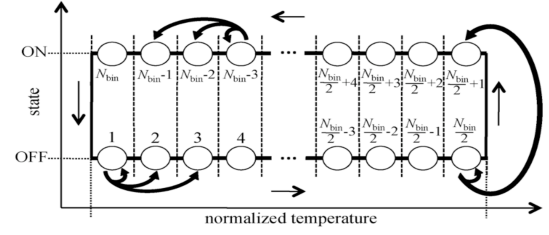


Figure 10: State Model for MKC-PS [16].

5.2 Comparison Scenario

The AC units in MB-RC implement the randomized controller described in (6). At the BA level the PI controller described in (15) is used with the gain values in Table 1, the same aggregate PI controller used with the CARC algorithm.

In MKC-PS [16] each AC's temperature is discretized and along with the on/off state used to create N_{bin} discrete classes. The state for a single AC is illustrated in Figure 10. At the aggregate level, the collection of ACs is modeled as a $N_{bin} \times N_{bin}$ finite state Markov chain, where the state is the fraction of ACs in each bin. An artificial control input $u \in \mathcal{R}^{N_{bin}/2}$, that represents the fraction of ACs that should switch on/off state in a on/off bin pair, is introduced.

In order to control the aggregate power consumption, a look-ahead controller is used at the BA level. This controller looks ahead one time step to determine how many AC's should switch mode state, termed u_{goal} , in order to meet the desired power consumption level. In order to convert this to the control input, u , u_{goal} needs to be distributed amongst the applicable bins. The authors of [16] offer two methods for this, (Controller 1) uniform distribution: each bin receives an equal amount, (Controller 2) preferential distribution: bins closer to the indoor temperature deadband receive preference. Once u_{goal} has been converted to u , the elements of u are divided by the corresponding state estimate to convert u to u_{rel} , a "broadcast probability vector". Upon receiving u_{rel} each AC can use its current state to determine which element of this vector to use as the probability of switching mode.

According to [16], the design of Controller 2 was to reduce AC short cycling. However in simulations we observed that Controller 1 in fact performed better in terms of cycling. We therefore use Controller 1 in our implementation for the simulation comparisons. Additionally, we give full state knowledge to MKC-PS, so the process of converting u to u_{rel} uses the true state (fraction of ACs in each temperature bin) and not the estimated state as described in [16]. Due to the above, we believe the simulation scenario and results presented for MKC-PS represent the best-case scenario.

5.3 Simulation Results

All results are for a closed loop system, where CARC and MB-RC implement a PI controller (15) at the aggregate level and MKC-PS uses the look-ahead controller, as described in [16]. For ease of interpretation the tracking results are converted from fraction of loads to MW. Figure 11 shows reference tracking for the CARC (Top) algorithm (our proposed design), MB-RC (Middle), and MKC-PS (Bottom). The results of the individuals cycling QoS are shown

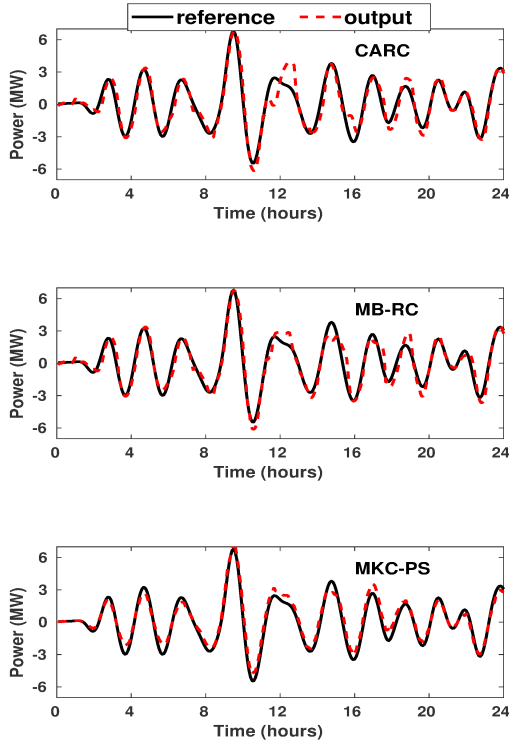


Figure 11: Tracking performance with reference r_k and output (power deviation), 10000 AC's. Top: CARC, Middle: MB-RC, Bottom: MKC-PS.

in Figures 12- 13 for all control scenarios. Additionally, applicable metrics for reference tracking and QoS are provided in Table 2.

Table 2: Simulation Results

Metrics	Controller		
	CARC	MKC-PS	MB-RC
tracking error (%)	4.4	3.6	4.1
Min Cycle Time (min.)	20	5	10
Max Device Switches	19	70	20

5.3.1 Global QoS: Reference Tracking Performance. The tracking results for all of the comparison control scenarios are shown in Figure 11. Numerically, the tracking results are tabulated in the first row of Table 2. We compute the tracking error in percentage as:

$$e_{track}^{(\%)} = \frac{100 \sum_{k=1}^{N_t} |(r_k + \bar{y}_k) - (y_k + \bar{y}_k)|}{\sum_{k=1}^{N_t} |(r_k + \bar{y}_k)|}$$

The three controllers perform similarly in regards to reference tracking percentage error.

5.3.2 Consumer QoS. There are two key individual QoS signals: indoor temperature and cycling of ACs. Recall that the temperature QoS is preserved through design: the local randomized controller

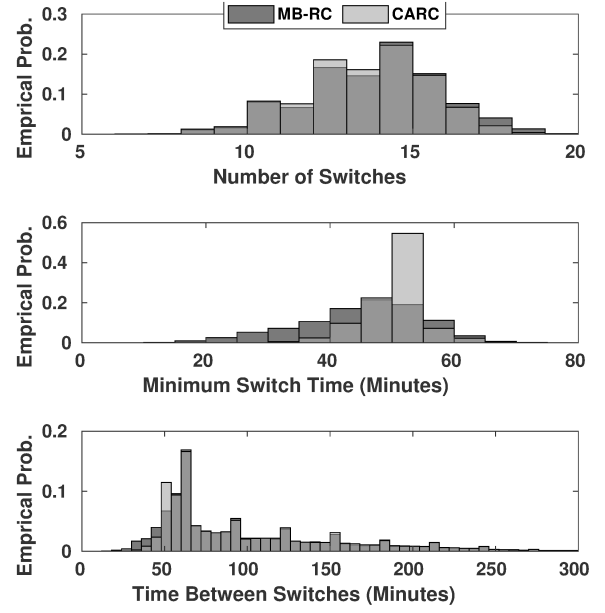


Figure 12: Empirical probabilities comparing CARC and MB-RC.

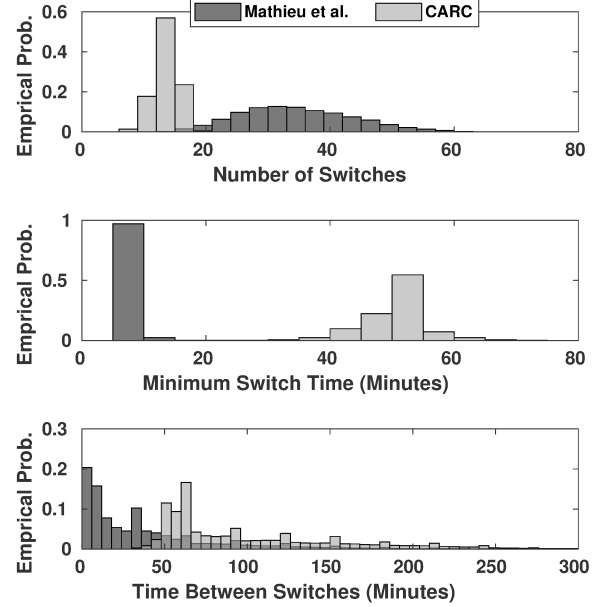


Figure 13: Empirical probabilities comparing CARC and MKC-PS.

is designed to behave identically to the deterministic thermostat controller when the temperature goes beyond the deadband.

To examine cycling QoS, defined in section 3.5 (CARC algorithm), we count the times between successive switches for an AC unit. The histogram of “time between switches for all AC’s is shown in

Figures 12 -13 (bottom). To test if the CARC algorithm has reduced short cycling, the minimum time between successive switches is computed for each AC (Figures 12 -13 (middle)). From these figures it is clear that the CARC Algorithm improves the minimum cycling time over both MB-RC and MKC-PS. In row 2 of Table 2, the minimum time between mode switches, searched over all ACs, is displayed. It is found that the CARC algorithm quadruples the minimum mode switching time when compared to MKC-PS and doubles when compared to MB-RC.

In addition to short cycling, the total number of switches an AC undergoes is investigated. The histogram illustrating total number of switches, computed over all ACs, is shown in Figures 12 -13 (Top). Between CARC and MB-RC the difference in the total number of switches is negligible. Although the AC with the greatest number of mode switches, under the CARC algorithm, switches once less than MB-RC (Table 2). Recall, that these experiments are performed over one day, so under similar daily conditions this would mean ≈ 30 less switches per month and ≈ 360 less switches a year. Comparing against MKC-PS the CARC algorithm drastically decreases the number of mode switches over all AC units.

6 CONCLUSION

Randomized control at the loads addresses many of the limitations of previous approaches to distributed coordination of TCLs, including issues of synchronization and computational complexity. It enables the BA to manipulate aggregate power consumption of a collection of ACs by broadcasting a single scalar to all ACs. The local randomized controller maintains the indoor temperature.

In this paper we leveraged the randomized control philosophy espoused in [3, 4, 7, 17] and extended it to also enforce the individuals cycling QoS constraint. We achieved this by adding local intelligence that forces a load to revert back to baseline behavior when it would otherwise cycle excessively because of the command from the BA.

The three control strategies, MKC-PS, CARC, and MB-RC, are all able to achieve similar reference tracking results. However, the CARC design is able to reduce cycling compared to MB-RC and MKC-PS design at little to no cost in tracking error.

A natural extension for future work is to incorporate a feed-forward component to the control system, so that errors in baseline prediction and reference signal generation can be systematically handled. The current control systems assumes that baseline power consumption and reference signal are known accurately.

Another path for future work is to propose a mechanism to preserve the individuals' privacy. The control architecture requires measurement of total power consumption of all the ACs taking part. Here we assumed that each AC sends on/off state information to the BA at every sampling period (5 minutes). It may be possible to infer patterns about individuals from this data that individuals may wish to keep private. A proposal in this direction is to randomly sample a small percentage of the total population to estimate the total power consumption [6], which shows good performance.

ACKNOWLEDGMENTS

The research reported here has been partially supported by the NSF through award 1646229 (CPS-ECCS) and ANR through award ANR-16-CE05-0008.

REFERENCES

- [1] Prabir Barooah, Ana Bušić, and Sean Meyn. 2015. Spectral decomposition of demand side flexibility for reliable ancillary service in a smart grid. In *48th Hawaii International Conference on Systems Science (HICSS)*. invited paper.
- [2] S. Bashash and H. K. Fathy. 2011. Modeling and control insights into demand-side energy management through setpoint control of thermostatic loads. In *Proceedings of the 2011 American Control Conference*. 4546–4553.
- [3] Ana Bušić, Md Umar Hashmi, and Sean Meyn. 2017. Distributed control of a fleet of batteries. In *American Control Conference (ACC)*. 3406–3411.
- [4] Ana Bušić and Sean Meyn. 2016. Distributed randomized control for demand dispatch. In *IEEE conference on decision and control*. 6964–6971.
- [5] Duncan S Callaway. 2009. Tapping the energy storage potential in electric loads to deliver load following and regulation, with application to wind energy. *Energy Conversion and Management* 50, 5 (2009), 1389–1400.
- [6] Yue Chen. 2016. *Markovian Demand Dispatch Design for Virtual Energy Storage to Support Renewable Energy Integration*. Ph.D. Dissertation. Gainesville, FL, USA. Advisor(s) Meyn, Sean P.
- [7] Yue Chen, Md Umar Hashmi, Joel Mathias, Ana Busic, and Sean Meyn. 2017. Distributed Control Design for Balancing the Grid Using Flexible Loads. In *IMA Volume on the Control of Energy Markets and Grids*. 1–26. <https://hal.archives-ouvertes.fr/hal-01656726>
- [8] Austin R. Coffman, Ana Busic, and Prabir Barooah. 2018. A Study of Virtual Energy Storage From Thermostatically Controlled Loads Under Time-Varying Weather Conditions. In *5th International High Performance Buildings Conference at Purdue*. 10.
- [9] Alexander W. Dowling and Victor M. Zavala. 2017. Economic opportunities for industrial systems from frequency regulation markets. *Computers & Chemical Engineering* (2017). <https://doi.org/10.1016/j.compchemeng.2017.09.018>
- [10] S. Kundu, N. Sinityn, S. Backhaus, and I. Hiskens. 2011. Modelling and control of thermostatically controlled loads. In *Power systems computation conference*.
- [11] Yashen Lin, Prabir Barooah, and Johanna Mathieu. 2017. Ancillary services through demand scheduling and control of commercial buildings. *IEEE Transactions on Power Systems* 32 (January 2017), 186 – 197. Issue 1. <https://doi.org/10.1109/TPWRS.2016.2557764>
- [12] Yashen Lin, Prabir Barooah, Sean Meyn, and Timothy Middelkoop. 2015. Experimental evaluation of frequency regulation from commercial building HVAC systems. *IEEE Transactions on Smart Grid* 6 (2015), 776 – 783. Issue 2.
- [13] M. Liu and Y. Shi. 2016. Model Predictive Control of Aggregated Heterogeneous Second-Order Thermostatically Controlled Loads for Ancillary Services. *IEEE Transactions on Power Systems* 31, 3 (May 2016), 1963–1971.
- [14] Y.V. Makarov, J. Ma, S. Lu, and T.B. Nguyen. 2008. Assessing the value of regulation resources based on their time response characteristics. *Pacific Northwest National Laboratory* (2008).
- [15] Johanna L. Mathieu, Mark Dyson, and Duncan S. Callaway. 2012. Using residential electric loads for fast demand response: The potential resource and revenues, the costs, and policy recommendations. In *Proceedings of the ACEEE Summer Study on Buildings*.
- [16] Johanna L. Mathieu, Stephan Koch, and Duncan S Callaway. 2013. State estimation and control of electric loads to manage real-time energy imbalance. *IEEE Transactions on Power Systems* 28, 1 (2013), 430–440.
- [17] Sean Meyn, Prabir Barooah, Ana Bušić, Yue Chen, and Jordan Ehren. 2015. Ancillary service to the grid from intelligent deferrable loads. *IEEE Trans. Automat. Control* 60 (March 2015), 2847 – 2862. Issue 1. <https://doi.org/10.1109/TAC.2015.2414772>
- [18] Cristian Perfumo, Ernesto Kofman, Julio H. Braslavsky, and John K. Ward. 2012. Load management: Model-based control of aggregate power for populations of thermostatically controlled loads. *Energy Conversion and Management* 55 (2012), 36 – 48.
- [19] Ana Radovanovic, William D. Heavlin, and Sila Kiliccote. 2016. Optimized Risk-Aware Nomination Strategy in Demand Response Markets. In *Proceedings of the 3rd ACM International Conference on Systems for Energy-Efficient Built Environments (BuildSys '16)*. ACM, New York, NY, USA, 99–108. <https://doi.org/10.1145/2993422.2993430>
- [20] W. Zhang, J. Lian, C. Y. Chang, and K. Kalsi. 2013. Aggregated Modeling and Control of Air Conditioning Loads for Demand Response. *IEEE Transactions on Power Systems* 28, 4 (Nov 2013), 4655–4664.